

1 Understanding Some Basic Statistical Concepts

1.1 Sample vs. Population

A concept that may not have been well explained in your statistics class of long ago, is sample vs. population estimates. If your mission is to measure weight on all patients with diabetes in the world (or what statisticians call the “universe”), we must measure the **mean** and **standard deviation** (SD) in a representative **sample** of patients that estimates the mean and variance of weight in the **universe** (μ or **standard error**, SE). Thus, we make estimates of population weight through a sample of diabetic patients, for example. By examining samples of data, we are always estimating characteristics of the universe. Ideally, these samples are randomly selected from the population and thus, if we were to choose another random sample from the same population, the results would be approximately equal, within the acceptable bounds of random error. This is why the sample selection becomes critical; the sample must be as representative of the population being studied as possible.

1.2 General Data Management Considerations

If you have a binary variable, i.e., a variable with only two choices, such as gender, in a dataset, it is necessary to always enter the binary code, or leave it blank if it is missing. Too often I have received datasets for analysis that contain a 1 for “yes” and blank for “no,” meaning to me that all of the blanks are missing. However, when I asked the author of the dataset about these missing values: “did the patient have cardiovascular disease (CVD)?,” she said “oh, no, a blank means that they did not have the CVD.” Those missing fields were corrected to receive a value of 0, while the truly missing

values were simply left blank. Make no assumptions if you see a lot of missing data.

Another tragic example was when a post-doc decided not to consult the codebook to define who was on study drug and who was on placebo, so he switched the assignments, and in his post-doctoral dissertation, he reported that patients on the study drug did not benefit while the controls did benefit, when indeed the reverse was true. Always document variable codes in a codebook or in the database itself.

Laboratory personnel also need to have access to statisticians, or they should possess a foundational and functional level of knowledge in statistics in order to understand, apply, and interpret their laboratory results and keep their instruments calibrated. Or, they are advised to speak to a statistician at the very start of study development. There can be struggles between laboratory personnel and statisticians/epidemiologists when it comes to data handling or interpretation. As one who helped a lab to optimize the performance of their assays, I would sometimes experience comments like, “oh, we eliminated the outliers,” or when data are missing for a certain field, they enter a QNS (quantity not sufficient) instead of leaving it blank, or entering dates as month/day/year in some of the fields and then day/month/year in others, thereby throwing off the date format recognition of my analysis software. So, my best advice is to speak to your statistician before designing your study and your database to discuss:

- study design
- developing the hypothesis and the null hypothesis
- sample size. . . sample size. . . sample size. . . sample size. . .
- Also, develop a data codebook (including strict formats for dates!).
- Keep ALL data and don't throw out the outliers!

1.3 Central Limit Theorem

Another related concept is the **Central Limit Theorem**, which simply posits that when one continually draws samples from a population and measures their HbA1c levels, for example, the HbA1c values from multiple subjects will eventually take on a normal distribution as one keeps sampling, that is, when plotted the values will take on a bell-shaped distribution that is centered around the mean and the median of the distribution, but only if

the values are capable of being normally distributed in the first place. For example, you wouldn't expect a logarithmically distributed variable to eventually take on a normal distribution after repeated sampling because it will always be logarithmically distributed. More on that in Section 2.2.

1.4 Parametric vs. Non-parametric Analyses

We will learn about different types of analyses to perform on different types of data, but the initial question to ask is: "Are the data normally distributed?" If yes, use the parametric statistics toolbox. If they are otherwise, i.e., not normally distributed, use the non-parametric toolbox.

Parametric statistics are a set of statistical procedures that are conducted on normally distributed, continuous variables. Parametric statistics are generally more robust than non-parametric statistics, so it becomes understandable why efforts are often made to "normalize" non-normally distributed data (see Chapter 4) before subjecting them to parametric statistics. For example, HIV viral load must be \log_{10} transformed to assimilate a normal distribution before being analyzed using the parametric *T*-test.

Non-parametric statistics are a set of statistical procedures that are performed on non-normally distributed data like binary, ordinal, and nominal variables, or on continuous variables that are not normally distributed.

Borrowed and adapted from Tanya Hoskin, a statistician in the Mayo Clinic Department of Health Sciences Research who provides consultations through the Mayo Clinic CTSA BERD Resource,¹ Table 1.1 elucidates which test to use in different circumstances by giving laboratory and clinical examples.

1.5 How to Calculate Some Basic Measures of Central Tendency

The mean, median, and mode are common indices used to describe the characteristics of a sample. They are simple to calculate and give some useful information on how sample values, like age and gender, are distributed; the indices can also be used to compare age and gender between different populations. However, the value of these indices has limitations, and misuse can yield misleading information. Following the definitions and the way to calculate the mean, median, and mode (below), two

Table 1.1 The selection of appropriate statistical tests is dependent on data type

Analysis type	Example	Parametric procedure	Non-parametric procedure
Compare means between two distinct and independent groups	Mean systolic blood pressure for patients on placebo vs. patients on study drug	Two-sample <i>T</i> -test	Wilcoxon rank-sum test
Compare two quantitative measurements taken from the same individual	Cell viability before vs. after 3 days in -80°F freezer	Paired <i>T</i> -test	Wilcoxon signed-rank test
Compare means between three or more distinct/independent groups	We want to compare the baseline ages of 3 groups: drug #1 vs. drug #2 vs. placebo	Analysis of variance (ANOVA)	Kruskal–Wallis test
Estimate the degree of association between two quantitative values	Viral particles in urine vs. saliva specimens	Pearson correlation	Spearman–Rank correlation

examples of populations are illustrated to make the point of the use and misuse of these measures.

1.5.1 Mean

The **mean** is calculated as the sum of the values divided by the number of the values. Why are means important? Because they give us a single summary value that describes one measure of the data that is useful for comparing across two or more populations.

Calculation of the mean:

What is the mean of 2, 3, 6, and 10?

Answer: $(2 + 3 + 6 + 10)/4 = 21/4 = 5.25$.

However, if you have a distribution such as the following:

What is the mean of $20 + 2 + 500 + 2500$?

Table 1.2 Frequency distribution of Measured Blood Loss

MBL				
		Frequency	Percent	Cumulative Percent
Valid	1.50	1	4.3	5.3
	1.54	1	4.3	10.5
	1.64	1	4.3	15.8
	1.78	1	4.3	21.1
	3.55	1	4.3	26.3
	3.70	1	4.3	31.6
	3.78	1	4.3	36.8
	4.27	1	4.3	42.1
	4.93	1	4.3	47.4
	6.41	1	4.3	52.6
	7.21	1	4.3	57.9
	7.39	1	4.3	63.2
	7.53	1	4.3	68.4
	7.84	1	4.3	73.7
	8.02	1	4.3	78.9
	8.63	1	4.3	84.2
	11.28	1	4.3	89.5
13.43	1	4.3	94.7	
68.83	1	4.3	100.0	
	Total	19	82.6	
Missing		4	17.4	
Total		23	100.0	

You might decide that finding the mean of this distribution may be meaningless since there is just so much space between the values. A way to reduce the space is by “normalizing” these values before taking the mean of them (see Section 1.4).

1.5.2 Median

The **median** is the midpoint of a frequency distribution where 50% of values fall below it and 50% of values fall above it. The median can be estimated by constructing a frequency distribution table.

As can be seen from Table 1.2, the midpoint of “Measured Blood Loss” can be found by looking at the cumulative frequency and finding that the

50% mark of the distribution falls somewhere between 4.93 and 6.41. These two values are almost equally distant from 50%, so we can approximate the median by taking the mean of these two values: $(4.93 + 6.41)/2 = 5.67 =$ the median.

1.5.3 Mode

The **mode** is the number that is repeated most frequently in a distribution. If there is a tie between two values, the distribution is said to be bimodal. If three or more values are tied, it is said to be multimodal. Mode can also be used in cases of ordinal variables like race; which race is most prevalent in the ordinal continuum of the race values?

Calculation of the mode:

What is the mode of 1, 4, 2, 4, 7, 5, 6, 4, 3, 7, 4, 4, 7, 7, and 7?

Answer: 4 and 7. This is a bimodal distribution.

Now, let's examine two different samples and see how the mean, median, and mode represent the samples and their usefulness.

Mean, Median, and Mode

Group #A: There are 200 individuals in a study that describes blood glucose levels. The data show that the blood glucose values seem to be normally distributed (see Introduction) from 50 to 150 mg/mL. That is, there are persons with low values, mid-values, and high values. If the data are plotted, they show a bell-shaped curve that is normally distributed (Figure 1.1).

Group #B: There are also 200 individuals from a different area and the distribution of blood glucose levels is assessed. In this group, it can be noted that about two-thirds of persons have very low glucose values (<50 mg/mL), while the upper third have very high levels (>150 mg/mL); there are very few persons with mid-range glucose values.

Interpretation of the mean and median in the two groups: when Groups A and B are combined, Population 1 appears to have HbA1c levels that are

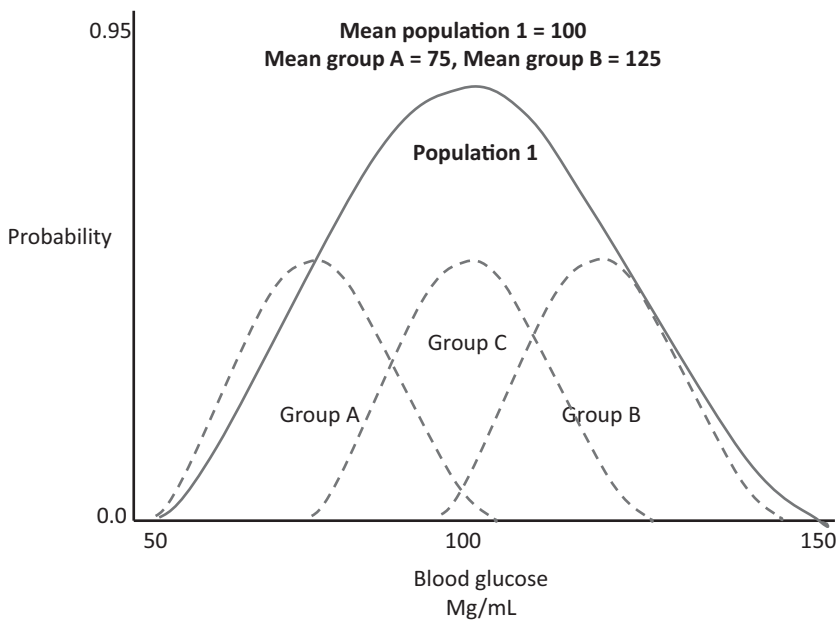


Figure 1.1 Dissimilar distributions of blood glucose levels.

normally distributed; the median is central and likely lies very close to the mean of 100. So, seeing that the population data are normally distributed, one may automatically consider running parametric analyses. However, when broken into Groups A and B and C, the group means of Groups A and B (i.e., 75 and 125), and likely their group medians, have shifted away from the population mean of 100 and the group distributions also appear to be statistically different from each other. Another visual observation is that Group C is likely not statistically different from Groups A and B. In this situation, a non-parametric statistic should be considered when the distributions of blood glucose in Groups A, B, or C are not normally distributed after being stratified from the population, even though the distribution of Population 1 appears to be normally distributed.

In summary, the characteristics of data are extremely important to understand, and therefore simple measures should not be used exclusively; other statistical tools must be considered to fully and correctly describe the data. These include the standard deviation, the coefficient of variation, the range, and the interquartile range.

1.5.4 Range

The lowest number and the highest number of a sorted distribution designates the **range** of the distribution. Ranges are useful when speaking of normal and abnormal ranges for a biological characteristic.

Calculation of the range:

Example for the Clinician

Using the Measured Blood Loss distribution in Table 1.2, the smallest number in the distribution of numbers is 1.50 and the largest number is 68.83. That is the range of values in the Measured Blood Loss distribution.

1.5.5 Interquartile Range

The **interquartile range** (IQR) is at the 25th and 75th percentile of the distribution. This is a useful set of numbers because it presents a little more information about how the data are distributed at a more granular level than the range. In normal situations, the 25th and 75th percentiles of distributions may not be conveniently obvious, as is shown in Figure 1.2,




Figure 1.2 Boxplot showing comparison of weight loss by gender.

where we must make decisions on where these distributional demarcations are, as shown in the example below.

Calculation of the IQR:

Example for the Clinician

Again using the distribution of Measured Blood Loss (Table 1.2), the 25th percentile lies somewhere between 1.78 and 3.55 (mean \sim 2.66). The 75th percentile is a judgment call in this case; since the cumulative percentage 73.9% is closer to the 75th percentile than is 78.3%, we can take the MBL value at the 73.9% value and take an MBL value of 8.02 to demarcate the 75th percentile of the distribution. Thus, 2.66–8.02 is the IQR of the distribution.

One of the more informative plots is the **boxplot** (Figure 1.2), which shows the 25th and 75th distributions, the means, the upper and lower bounds of the 95% confidence intervals, and the points outside the 95% confidence interval, which are otherwise known as “outliers.” Boxplots are particularly informational and are attractive representations for illustrating statistical differences in journal articles.

Evidently, males were not significantly heavier than females at 24 months post-surgery, as can be determined by the overlap in confidence intervals (Section 3.2).

Knowing these basic statistical calculations and permutations can lead to a broader understanding of your data in terms of their distributions. This understanding is a crucial element in choosing the correct statistical approach. Know your data!

1.5.6 Skewness

Non-normally distributed data do not resemble a bell-shaped curve but rather, they may be skewed to the left or to the right as shown in Figure 1.3. **Skewness** is a term referring to the tail of the distribution, whether it leads off to the right or the left. The skewness of a data distribution refers to the symmetry of the values around the mean of the distribution.

Why do we care about skewness? Because it shows us if the data are normally distributed and therefore, when we should and should not use




Figure 1.3 Depiction of negative and positive skewness.

parametric analysis techniques. When the mean is equal to the median (is normally distributed), the value for skewness is 0, and is therefore appropriate for parametric analyses. When the majority of values in your distribution are concentrated together at the right or the left, the tail falls off to the right or the left and is therefore called “skewed to the right” or “skewed to the left,” respectively, as shown in Figure 1.3, and it is not appropriate for parametric analyses. Skewness values between -2 and 2 are generally acceptable for demonstrating that data are normally distributed, but values less than -2 and greater than 2 are good indications of skewness, or they are not normally distributed, and again, are not good candidates for parametric analyses. For an online skewness calculator, you can enter your data in a tool.²

Notice that when data are positively skewed, the mode and the median lie to the left of the mean, while in negatively skewed data, the median and the mode lie to the right of the mean.

The following is a perfect case in point. Presented in Figure 1.4(a) is a frequency distribution and in Figure 1.4(b) a bar chart of a continuous variable: “Days of Hospital Stay.” Upon visual scrutiny, this continuous variable is skewed to the left, since most cases stayed 0 or 1 day; thus, Days of Hospital Stay is not normally distributed around the mean. Also, in the cumulative frequency column of the table, you can see that 79.9% of the sample had a length of hospital stay of 2 days or less. Now, had you assumed that Length of Hospital Stay was normally distributed and you ran it through a parametric procedure, your result would be erroneous

(a)

Days of Hospital Stay	Frequency	Percent	Valid percent	Cumulative percent
0	3097	36.9	44.5	44.6
1	1475	17.6	21.2	65.8
2	979	11.7	14.1	79.9
3	390	4.6	5.6	85.5
4	209	2.5	3.0	88.5
5	151	1.8	2.2	90.7
6	123	1.5	1.8	92.4
7	102	1.2	1.5	93.9
8	62	0.7	0.9	94.8
9	60	0.7	0.9	95.6
10	42	0.5	0.6	96.3
11	30	0.4	0.4	96.7
12	24	0.3	0.3	97.0
13	28	0.3	0.4	97.4

(b)

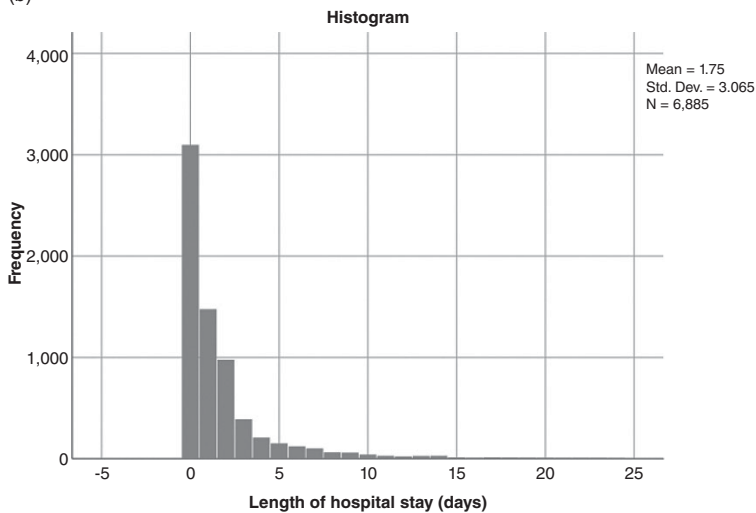


Figure 1.4 (a) Frequency distribution and (b) bar chart of non-normally distributed data.

since it neglects the assumption of distributional normality where skew is equal to zero. More information is available on normality.³ As an alternative to transforming the data to assume a normal distribution, it is advisable to use non-parametric techniques (Chapter 5).

So, by entering your data into the calculator,⁴ the skewness is 2.45, which exceeds the 2 boundary and is thus a candidate for non-parametric analyses, since it would not be possible to transform the data as they do not take on any sort of predictable distribution, such as a logarithmic distribution (Section 2.2).

1.5.7 Kurtosis

Again, examining the distribution of your data before assuming normality is essential when considering whether or not to perform parametric or non-parametric techniques for analysis. Measuring the **kurtosis** of your data is another way to examine the shape of the distribution of your data. As shown in Figure 1.5, kurtosis is specifically a measure of the tails of your distribution. You can see that the tails intercept the x -axis rather sharply (where $\sigma = 1$) or more gently (where $\sigma = 3$). If $\sigma = 3$, your data are normally distributed. If σ is greater than 3 or less than 3, your data are said to be kurtotic, i.e., non-normal. Notice that when you run your data through the




Figure 1.5 Measures of kurtosis.

skewness calculator,⁵ it also generates a value for kurtosis. You may use these values to decide if your data are or are not normally distributed and apply the appropriate set of statistics accordingly, i.e., parametric or non-parametric.

1.6 Frequency Distributions

In Table 1.2 we examined the distribution of Measured Blood Loss during surgery. This is called a **frequency distribution** of all data points in my data file. When a statistician receives a new dataset, the first exploratory analysis they often run is a frequency distribution as a first step to check the data for accuracy and coding errors. They may run “frequencies” on all variables in the dataset. The utility of the frequency distribution is that one can easily identify values that are out of range and values that are character instead of numeric, or vice versa. One may also run frequency distributions to find meaningful cutoffs if one wishes to categorize these variables by the mean, the median, or the 25th percentiles – e.g., age (0–2, 3–5, 6–12), weight (90–100, 100–150, 150–200), etc. These are common data transformations that are made to continuous values before entering the newly coded variables into other analyses.

For example, see Table 1.3.

Table 1.3 Utility of frequency distributions

	Race	Frequency	Percent	Valid percent	Cumulative percent
Valid	0.0	5	0.1	0.1	0.1
	1.0 Caucasian	4656	67.7	68.0	68.0
	2.0 African American	1854	27.0	27.1	95.1
	3.0 Asian	71	1.0	1.0	96.1
	4.0 Hispanic	161	2.3	2.4	98.5
	5.0 Other	101	1.5	1.5	100.0
	6.0	2	0.0	0.0	100.0
	9.0	1	0.0	0.0	100.0
	Total	6851	99.6	100.0	
	Missing	System	25	0.4	
Total		6876	100.0		

By running a frequency distribution, you can note in Table 1.3 that the 0s, 6s, and 9s are errors and that race is missing for 25 records in your raw data. This simple table should prompt the data crew to find these missing values and complete them, and also find the records containing the errors and fix them. This process is known as “cleaning your data.”

Frequency Distribution

Moving from left to right in Table 1.4, the value in the first column is the actual value for number of post-surgical complications.

Table 1.4 How to interpret a frequency distribution table

Number of post-surgical complications	Frequency	Percent	Valid percent	Cumulative percent
0	6420	76.4	97.7	97.7
1	131	1.6	2.0	99.7
2	16	0.2	0.2	100.0
3	1	0.0	0.0	100.0
Total	6568	78.2	100.0	
Missing	1832	21.8		
Total	8400	100.0		

The **frequency** represents the number of times the value appears in the dataset. The **percent** column shows the proportion of that value within the whole distribution, including those with missing values. There are 8400 total records in the database, of which 1832 are missing values for number of post-surgical complications. The **valid percent** shows the proportion of that value in your database when you *exclude* missing values. And finally, the **cumulative percent** shows the increasing (or cumulative) percentage of values from the lowest value in the distribution to the highest value (0% to 100%). When reporting frequencies in manuscripts or reports, it is best to state the total number of records in the database, the number of missing values, and the valid percentage of values in the dataset.

Because Estimated Blood Loss (EBL) is a continuous value (Table 1.5), one would tend to examine the mean and standard deviation of EBL,

Table 1.5 A frequency distribution of Estimated Blood Loss

Estimated Blood Loss (in dL)	Frequency	Percent	Valid percent	Cumulative percent
100.00	1	4.3	4.3	4.3
150.00	3	13.0	13.0	17.4
200.00	1	4.3	4.3	21.7
300.00	4	17.4	17.4	39.1
400.00	3	13.0	13.0	52.2
450.00	1	4.3	4.3	56.5
500.00	4	17.4	17.4	73.9
800.00	2	8.7	8.7	82.6
1000.00	2	8.7	8.7	91.3
1100.00	1	4.3	4.3	95.7
6250.00	1	4.3	4.3	100.0
Total	23	100.0	100.0	

Table 1.6 Making clinically relevant stratifications of continuous data

EBL (dL)	Mean baseline weight	SD	<i>P</i>
≤200	214.53	10.22	
201–400	220.64	12.32	<0.001
401–500	228.43	10.56	
≥501	258.22	13.65	

which is an appropriate way to summarize normally distributed, continuous data.

For example, look at mean weight compared across categories of EBL groupings of the EBL distribution (Table 1.6). When grouped into such clinically meaningful quartiles, i.e., ≤200, 201–400, 401–500, 501+, you can then use these categories to compare another continuous variable, say, the mean baseline weight between these EBL categories to answer the question: Do patients who lose more blood during surgery weigh more than those who lose less blood?

So, the importance of this example is to show the utility of frequency distributions and how they can aid in the manipulative grouping of data in order to conduct meaningful statistical analyses.

1.7 Measures of Dispersion and Variance

1.7.1 Standard Deviation

The standard deviation is one way to describe the variation of your data in relation to the mean of your data. It is a sum of the difference between each data point and the sample mean divided by the number of records in your dataset. The distance of each value from the mean is squared to ensure we do not obtain negative values, since the standard deviation must always be an integer whether it be lower or higher than the mean. Later, we take the square root of this difference divided by n , to negate the squaring:

$$\sigma = \sqrt{\frac{\sum(x - \bar{x})^2}{n}}$$

where

σ = standard deviation

\sum = sum of

x = each value in the dataset

\bar{x} = mean of all values in the dataset

n = number of values in the dataset.

Calculation of the standard deviation:

What is the standard deviation of 2, 3, 6, and 10?

Answer: Firstly, find the mean of these four values:

$$2 + 3 + 6 + 10 = 21$$

$$21/4 = 5.25$$

$$\begin{aligned} \text{SD} &= \sqrt{\frac{(2 - 5.25)^2 + (3 - 5.25)^2 + (6 - 5.25)^2 + (10 - 5.25)^2}{4}} \\ &= \sqrt{\frac{10.56 + 5.06 + 0.56 + 22.56}{4}} \\ &= \sqrt{9.685} = 3.11. \end{aligned}$$

It might be more intuitive to make the calculations of the standard deviation in tabular form, as shown in the next example.

Standard Deviation

Table 1.7 shows the temperature in a very hot African city over a 20-day period and the 20-day average (mean) temperature for the 20 consecutive days.

Table 1.7 Calculation of standard deviation

	Close	20-Day mean	Deviation	Deviation squared
1	109.00	112.30	-3.30	10.91
2	103.06	112.30	-9.24	85.38
3	102.75	112.30	-9.55	91.26
4	108.00	112.30	-4.30	18.52
5	107.56	112.30	-4.74	22.47
6	105.25	112.30	-7.05	49.75
7	107.69	112.30	-4.62	21.30
8	108.63	112.30	-3.68	13.53
9	107.00	112.30	-5.30	28.12
10	109.00	112.30	-3.30	10.91
11	110.00	112.30	-2.30	5.30
12	112.75	112.30	0.45	0.20
13	113.50	112.30	1.20	1.43
14	114.25	112.30	1.95	3.79
15	115.25	112.30	2.95	8.68
16	121.50	112.30	9.20	84.58
17	126.88	112.30	14.57	212.34
18	122.50	112.30	10.20	103.97
19	119.00	112.30	6.70	44.85
20	122.50	112.30	10.20	103.97
	2246.06	112.30		921.28
				DevSqr/20 46.06
				StdDev 6.787

To calculate the standard deviation, the distance of each daily temperature from the mean temperature is calculated and then “squared” to obtain a positive value for the negative deviations from the mean. These squared distances are then summed together and divided by the n of your sample. Then, the square root of that value is taken to normalize the squared value back to the unsquared value (the correction).

1.7.2 Standard Error

The standard error is a value that pertains to the variance of the population mean, not the sample mean! Mathematically, it is SD/\sqrt{n} . Conceptually, it is a measure of how precise a parameter is in the population. That parameter can be the mean or the correlation coefficient (Section 4.3.3). If we were to repeatedly pull samples from the same population and measure the mean weight for each sample, the SE of all the repeated weight measurements is the measure of precision of all sample mean estimates.

1.7.3 Coefficient of Variation

The coefficient of variation (CV) is another measure that describes the dispersion of measurements in terms of the standard deviation from the mean. It is commonly used in reproducibility and repeatability studies to determine how steady the measurements are (or are not) from an identical sample or person. Usually, if you are repeatedly testing the same sample over and over, it is preferred that the value for the CV is small. The CV is calculated by dividing the standard deviation by the mean of a distribution which is derived from the same sample or subject. It is just an expression used to show the ratio of the standard deviation from the mean.

CV for a sample:

$$CV = \frac{SD}{\bar{x}} \times 100\%$$

where SD is the standard deviation and \bar{x} is the sample mean.

Using the values above, the mean (5.25) and standard deviation (3.11) give a CV of 0.59, or 59%, which is quite high. Keep in mind that this calculation of CV is only appropriate for normally distributed measurement data, meaning that one would not normally calculate a mean or standard deviation or CV for heavily skewed data (Section 1.5.6).

Calculation of the coefficient of variation:

Example for the Laboratorian

Laboratory A and Laboratory B have each been given 10 replicate aliquots of blood from *one healthy volunteer* to test for blood glucose. Both labs performed their tests using the same measurement

procedure, used the same tester, the same measuring instrument under the same conditions, in the same location, and the test repetition occurred over a short period of time.

Laboratory A gets a mean of 81.1 and SD of 15.7.

Laboratory B gets a mean of 82.5 and SD of 20.4.

Which lab had better test repeatability?

Laboratory A $CV = 15.7/81.1 \times 100 = 19.35\%$.

Laboratory B $CV = 20.4/82.5 \times 100 = 24.73\%$.

Clearly, Laboratory A had less variation in blood glucose levels in relation to the mean value, and thus had better test repeatability.

CVs are also computed when reproducibility or precision (Section 8.3.3) of one test is being assessed for variation. Ideally, the CV of a test should not change if tested in two different labs; that is, the one test should perform identically in both labs.

Links to Online Tools

- ¹ www.mayo.edu/research/documents/parametric-and-nonparametric-demystifying-the-terms/doc-20408960
- ² www.socscistatistics.com/tests/skewness/default.aspx
- ³ www.wessa.net/rwasp_skewness_kurtosis.wasp#output
- ⁴ www.socscistatistics.com/tests/skewness/default.aspx
- ⁵ www.socscistatistics.com/tests/skewness/default.aspx